

On the neural computation of utility

Peter Shizgal and Kent Conover

Center for Studies in Behavioural Neurobiology

Department of Psychology

Concordia University

Montréal, Québec, Canada H3G 1M8

This article appeared originally in *Current Directions in Psychological Science*, 1996, **5**(2), 37-43 (© American Psychological Society) and is reprinted with the permission of Cambridge University Press. All reproduction, copying or distribution of this material in any format, beyond single copying by an authorized individual for personal use, must first receive the written consent of Cambridge University Press (www.cup.org).

This work was supported by a grant (#ER0124) to Peter Shizgal (p.i.), Shimon Amir, and Barbara Woodside from the "Fonds pour la Formation de Chercheurs et l'Aide à la Recherche du Québec" and a grant (#A0308) to Peter Shizgal from the Natural Sciences and Engineering Research Council of Canada (NSERC). Kent Conover was supported by post-doctoral fellowships from the "Fonds de la Recherche en Santé du Québec" and from NSERC.

The authors thank Randy Gallistel and Alan Spector for their helpful comments and Shimon Amir, Steve Cabilio, and Luigi Riscaldino for their assistance in preparing the figures.

The self-stimulating rat presents a compelling spectacle. Having been trained to press a lever that triggers intense, continuously available stimulation of a site in the medial forebrain bundle, the rat works in a frenzied, insatiable fashion, even at the cost of forgoing its sole daily opportunity to obtain food. The ardor and determination shown by the rat suggest that obtaining additional stimulation has become an extraordinarily important goal. That this should be so is perplexing. If one were to insert a test probe into the central processing unit of a computer and deliver trains of current pulses, one would hardly expect to inject meaningful data. How could a signal meaningful to a rat arise from delivery of synchronous stimulation via a stout wire crudely inserted into the intricate fabric of the brain? If the induced neural activity is somehow meaningful, what natural signal does it mimic?

On the basis of experiments on the relationship between the rewarding effects of electrical brain stimulation and gustatory stimuli^{1,2,3}, we have proposed a new account of the nature of the electrically evoked signal. In this essay, we flesh out our account by considering the phenomenon of brain stimulation reward (BSR) in relation to the computational processes involved in goal selection. By so doing, we address the function of the underlying neural circuitry and the question of how the electrical stimulation produces an apparently meaningful effect.

Central to our formulation is the concept of utility, which we have borrowed from economics. We assume our rats to be rational consumers insofar as they will prefer, under non-satiating conditions, an alternative that provides more of a given appetitive goal object (e.g., food) over an alternative that provides less. The relative utility of two different goal objects will depend not only on their abundance but also on the physiological state of the consumer and the ecological context in which the goal objects are embedded. In effect, we treat utility as a subjective estimate of the potential contribution of a goal object to fitness. The more accurate the estimate, the more adaptive are the choices that take the utility value into account.

In natural settings, the goals competing for behavior are complex, multidimensional objects and outcomes. Yet, for orderly choice to be possible, the utility of all competing resources must be represented on a single, common dimension⁴. In our view, BSR arises from the electrical activation of neurons that implement a unidimensional representation of the utility of natural goal objects, and it is the unidimensional character of information coding in this population of neurons that enables the electrical stimulation to produce a meaningful signal. In contrast, we argue, processing must be multidimensional at the earlier stage of processing where physiological feedback exerts its specific influence on goal selection. Thus, we see BSR as similar to natural rewards in terms of how utility is computed and used to select goals. However, we see BSR as quite different from natural rewards in terms of the modulating effect of physiological feedback and the sensory processing that accompanies the computation of utility. Although the rat in our portrayal does not hallucinate a piece of cheese upon receiving the electrical stimulation, the rat employs common neural circuitry in its evaluation of electrical and gustatory rewards.

Experimental Results

To illustrate these ideas, we summarize the results of several recent experiments^{1,2,3} and discuss their implications. The purpose of these experiments was to examine the relationship between BSR and rewarding effects of natural stimuli. Figure 1a illustrates the experimental preparation and one of the testing paradigms. In addition to a stimulation electrode aimed at the lateral hypothalamic (LH) level of the medial forebrain bundle, the rat is equipped with an intraoral catheter and an intragastric cannula. Brain stimulation is an unusual reward in that a single response suffices both to procure and “consume” it. The presence of an intraoral catheter connected to an infusion pump confers this property on a gustatory reward, a 55 to 85 μ l infusion of a highly concentrated (1 M) sucrose solution. By touching an empty drinking spout, the rat triggers the pump, which infuses the sucrose directly into the rat’s mouth. Touching the second

of the two empty spouts triggers the delivery of a train of electrical stimulation pulses. A second unusual property of BSR is the absence of satiation. The intragastric cannula renders the gustatory reward similar to BSR in this respect; as the rat swallows the sucrose solution, it drips out of a drain tube attached to the gastric cannula, thus minimizing postingestive effects.

The rat's task in the first experiment was to choose between the two rewards by touching the appropriate spout. After each choice, the spouts were disarmed, and a 2-s interval elapsed before another choice was made available. The number of each kind of reward earned was recorded at the end of each 2-min trial. The **standard** (sucrose) reward was held constant across trials, whereas the strength (the number of pulses per train) of the **alternate** (brain stimulation) reward was varied across trials. Additional test sessions were run with the pump turned off and only the electrical reward available.

Figure 1b shows results obtained from 1 subject: The number of choices of each reward is shown as a function of the number of pulses in the alternate (BSR) reward. The solid symbols represent the results obtained when the two rewards were in competition: Circles represent the number of sucrose rewards earned per trial, and triangles represent the number of electrical rewards earned. When the rewarding trains of electrical stimulation were composed of relatively few pulses, the rat chose the sucrose reward almost exclusively. Increasing the number of pulses reversed the preference. Note the rightward displacement of the filled triangles from the open triangles. The filled triangles represent choices of the BSR when it competed with the sucrose, whereas the open triangles represent choices of the BSR when the sucrose was unavailable. The presence of the sucrose led the rat to forgo moderate-strength trains (34 and 38 pulses) of brain stimulation for which it would have worked had the sucrose been unavailable. This would not have occurred had the rat used a categorical rule, selecting the brain stimulation whenever its value was above a certain threshold. Such categorical choice does not require a common evaluation of the two rewards. In contrast, the results imply that on a given trial, the rat selected the alternative that registered a larger value in a common system of measurement.

In a subsequent experiment, we determined whether our subjects could combine the values they assigned to the electrical and gustatory rewards. The experimental protocol is summarized in Figure 2a. In this experiment, the standard reward (held constant across trials) consisted of both an intraoral infusion of sucrose *and* an equally preferred train of brain stimulation (43 pulses per train, the intersection of the fitted functions in Fig. 1b); the alternate reward consisted of a train of brain stimulation, which varied in strength from trial to trial. Sample results are presented in Figure 2b. The filled circles represent the number of compound (sucrose + stimulation) rewards earned per trial, whereas the filled triangles represent the number of electrical rewards earned. Heavy lines (“Sum”) represent broken-line functions, consisting of lower and upper horizontal segments joined by a linearly increasing segment, fitted to the number of rewards earned at each value of the brain stimulation alternate. The fine lines (“Comp”) represent broken-line functions fitted to the results of the competition experiment (Fig. 1b). The displacement of the heavy lines from the fine ones indicates that the rat assigned a higher value to the compound reward than to its sucrose component alone. When brain stimulation was added to the sucrose, the strength of the alternate BSR had to be increased by 20 pulses (from 48 to 68) in order for the rat to forgo the compound standard. Thus, reward summation occurs when a small spritz of sucrose into the mouth is accompanied by a brief zap delivered to the LH.

Like the results in Figures 1b, the results in Figure 2b imply that the electrical stimulation and the sucrose were subjected to a common evaluation. Note that summation between two things is impossible unless they share the property assessed by the system of measurement. In principle, we could use a pan balance, but not a voltmeter, to measure summation between neutrons and electrons because both particles possess mass whereas only one of them possesses net charge. By analogy, summation between LH stimulation and intraoral sucrose is manifested in our experiment because both possess a common property: the ability to serve as a reward for operant performance. We speculate that this common ability arises from a common action of the

gustatory and electrical stimuli on a neural system that determines goal selection by signaling the utility of competing goals.

We performed two additional experiments to shed light on the locus of action of the electrical stimulation within the evaluative circuitry. First, we increased the value of a gustatory stimulus by creating a physiological need for it. Second, we decreased its value by allowing an ingested solution to accumulate in the gut. If, as some researchers have proposed, gustatory and electrical rewards are affected similarly by changes in physiological state, then the manipulations we performed should not have altered their relative values. However, if the rewarding effect of brain stimulation is combined with the rewarding effect of tastants downstream from the point or points where signals reflecting physiological state adjust the value of gustatory stimuli, then the manipulations should have left the value of the brain stimulation unchanged while altering the value of the gustatory stimuli.

In the first of the experiments on the locus of action of the electrical stimulation², we depleted the rats of sodium by administering a diuretic. Figure 1c depicts competition between a 0.9% saline solution and LH stimulation in a sodium-depleted rat. Note that as in Figure 1b, the rat gave up suprathreshold trains of LH stimulation in order to earn gustatory rewards: The curve relating the number of brain stimulation rewards earned to the strength of the stimulation (the reward-number curve) was displaced to the right when the electrical stimulation competed with the gustatory reward. This rightward shift was absent when the subject was sodium replete (not shown). In contrast to the large effect of sodium depletion on the value of the gustatory reward, the value of the electrical reward appeared essentially unchanged. As shown in Figure 1d, similar reward-number curves were obtained for the BSR alone in the sodium-depleted (open triangles) and sodium-replete (filled triangles) states.

In the second experiment on the locus of action of the electrical stimulation³, we assessed the impact of postingestive feedback, modifying the summation paradigm depicted in Figure 2a so

that the strength of both the alternate (electrical stimulation alone) and the standard (sucrose + stimulation) reward were held constant from trial to trial (Fig. 3a). When the gastric cannula was open (results for 1 subject are shown in Fig. 3c), preference was stable over the 30-min test session. In contrast, when the cannula was closed (Fig. 3d), preference for the compound reward was abolished by the end of the test session. In 2 other subjects (results not shown), the preference reversed during the test session so that by the end, the rats chose the stimulation alone over the combination of the same stimulation train and an intraoral infusion of sucrose.³ In contrast to the dramatic effect of closing the gastric cannula on preference between the electrical stimulation and the compound reward, closing the cannula failed to alter rate-number curves for the stimulation alone. Figure 3b shows very similar rate-number curves obtained before and after test sessions conducted with the cannula closed (Fig. 3d), sessions in which the cumulative intake of sucrose averaged 30 ml in 30 min (a very large meal). These results and those of an additional control experiment imply that postingestive feedback from prodigious self-administered gastric loads undermine the value of the gustatory reward without altering substantially the value of the electrical reward. In some subjects, the gustatory stimulus became aversive by the end of the test session. Thus, the change in physiological state caused the value of the gustatory stimulus, to vary in both magnitude and sign, as in Cabanac's⁵ demonstrations, whereas the value of the electrical reward was perturbed little, if at all.

Discussion

With sample results now before us, let us reconsider the formulation presented at the beginning of this essay.

Currency functions: Imagine a shopper confronted with a series of tasks of increasing difficulty. First, the shopper must decide whether the potatoes offered by one grocer for \$2.00 per kg represent a better deal than apparently-identical potatoes offered by another grocer for 10 French francs per kilogram. The second task is to decide whether a rich cheese at \$10.00 per kg is a

better buy than lean ham at \$10.00 per kg. Finally, we ask the shopper to determine which of two different “packages” of goods offers the better value, a \$2.50 sandwich made from the rich cheese alone or a \$3.00 sandwich containing both cheese and ham.

The first task illustrates why adaptive choice requires ranking of alternatives along a single dimension. As long as two different currencies (scales of value) are used, the consumer cannot decide which grocer offers the best deal. The problem is solved by converting both prices to a common currency. The second task is more difficult because it involves comparison of different goods; to obtain a maximal payoff, our shopper must compute the relative utilities of the cheese and ham at the moment of choice. We presume that this computation is influenced by factors such as past experience and current physiological state. For example, choice might be biased toward the cheese if the consumer had endured a recent case of gastric distress after eating ham or had recently undergone a large reduction of fat stores. The third task illustrates a combinatorial problem: determining the combined utility of ham and cheese in the same units as the utility of cheese alone.

For animals foraging in the natural environment, it is the third problem that is the most realistic. Different foodstuffs are distributed differentially in the environment. Each foodstuff is a chemically complex package of resources germane to the operation of different regulatory systems such as those controlling energy, mineral, and fluid balance. If the animal is to maximize return on its foraging efforts, it must weight the components of each food by past experience and by the state of the relevant physiological variables and then combine the weighted values so as to obtain an overall assessment of the utility of each complex foodstuff.

The competition results in Figures 1b and 1c depict the solution of a problem analogous to the second task. The rat chose between two different goods, brain stimulation and a tastant. The cost of the two items was same (both were available on the same schedule of reinforcement), and the amount of brain stimulation available for this cost was varied from trial to trial. The rat

manifested orderly choice between the two different goods, apparently choosing the higher valued item. Hence, we propose that the signal produced by the stimulation is evaluated in the same currency as the gustatory reward.

The summation test described in Figure 2 is roughly analogous to the third problem. The rat chose between one package containing a single commodity, brain stimulation, and another containing two commodities, brain stimulation and a tastant. Like the hypothetical consumer, the rat appears to combine the results of evaluating the different commodities in a package, thus enabling it to choose adaptively between packages.

Implications for the computation of utility: The results in Figures 1c, 1d, 3b, 3c, and 3d suggest that the utility calculation took physiological state into account in the case of the tastant but not in the case of the brain stimulation. Two questions arise from this conclusion. First, at what stage in the calculation of utility can physiological state exert a specific influence? As summarized in Figure 4, we argue that at the stage where physiological feedback specifically alters utility, the representation of the stimulus must be multidimensional, retaining both qualitative and quantitative information. Presumably, signals reflecting energy balance modulate the value of sucrose, whereas signals reflecting mineral and fluid homeostasis modulate the value of saline; the stronger the imbalance, the stronger the modulation. If these signals acted at a stage of evaluation where tastants were represented only by single values, then the effect of the physiological signals could not be confined to stimuli of a particular kind. Negative energy balance would increase the value of both sucrose and saline, as would negative sodium balance. However, if the physiological variables were incorporated in the computation of utility at a stage where the sensory quality (sweet vs. salty) of the tastant were preserved, then energy balance and sodium balance could exert categorically different effects, thus marshaling behavioral choice to contribute to physiological homeostasis. In future experiments, it will be important to verify

whether sodium depletion and depletion of energy stores indeed exert such specific influences on gustatory evaluation.

A second question arises from the results in these figures: why were sodium depletion and postingestive feedback ineffective in modulating the value of BSR? We think these findings imply that the utilities of brain stimulation and the tastants are combined downstream from the point where postingestive feedback modulates the value of the tastants. (See Fig. 4.) We propose that at the stage of the evaluative circuitry where the brain stimulation acts, the multiple dimensions of a natural stimulus have been collapsed into one, and the signal produced by sucrose is of the same kind as the signal produced by saline. The information about physiological state has already played its role and does not contribute further to the computation of the relative utility of different goal objects.

Why the effect of the stimulation is meaningful. The notion that processing is unidimensional in the stage of the evaluative circuitry activated by the brain stimulation offers an answer to the question of how the artificially patterned stimulation can elicit a meaningful effect. We will develop this argument by analogy to information processing in sensory systems. Sensory nerves carry signals that encode many different dimensions of a stimulus; hence, for example, gross stimulation of the auditory nerve evokes noisy, multitonal sensations⁶. In contrast, activity in a population of adjacent cells in sensory cortex may represent a point along a single stimulus dimension; hence, for example, stimulation of a localized region of V5 (a region where the direction of visual motion is represented) has systematic effects on the motion judgments of a monkey, as shown by Salzman, Britten, and Newsome⁷. We would argue that electrical stimulation in the experiment of Salzman et al. produced a meaningful signal because coding was unidimensional in the population of stimulated cells: This would be so, for example, if it were the aggregate activity in the stimulated neurons, and not its precise spatiotemporal distribution, that influenced the perceived direction of motion. Analogously, we would argue that electrical

stimulation of the LH produces a meaningful signal because this stimulation too activates a neural system at a unidimensional stage of processing. Indeed, there is already strong evidence that the reward value of LH stimulation is determined by the aggregate activity evoked (total impulse flow in the stimulated substrate) and not by the spatiotemporal distribution of the evoked activity: The same reward value can be produced by activating a large number of reward-related neurons at low frequency or a small number at high frequency⁸.

Future directions. With the aid of psychophysical, electrophysiological, immunocytochemical, and lesion methods, we and our co-workers are striving to identify the directly activated neurons subserving BSR. By recording from these cells in subjects performing decision tasks, by describing the neural network that furnishes these cells information about the external and internal milieus, and by further behavioral study of the evaluation of BSR and natural goal objects, we hope to gain insight into how the computation of utility is implemented in the brain.

References

1. K. L. Conover and P. Shizgal, Competition and summation between rewarding effects of sucrose and lateral hypothalamic stimulation in the rat, *Behavioral Neuroscience*, **108**(3), 537-548 (1994).
2. K. L. Conover, B. Woodside, and P. Shizgal, Effects of sodium depletion on competition and summation between rewarding effects of salt and lateral hypothalamic stimulation in the rat, *Behavioral Neuroscience*, **108**(3), 549-558 (1994).
3. K. L. Conover and P. Shizgal, Differential effects of postingestive feedback on the reward value of sucrose and lateral hypothalamic stimulation in the rat, *Behavioral Neuroscience*, **108**(3), 559-572 (1994).
4. D. J. McFarland, and R. M. Sibley, The behavioural final common path, *Philosophical Transactions of the Royal Society of London (Series B)*, **270**, 265-293 (1975).
5. M. Cabanac, Physiological role of pleasure. *Science*, **173**, 1103-1107 (1971).
6. G. E. Loeb, The functional replacement of the ear, *Scientific American*, **252**(2), 104-111 (1985).
7. C. D. Salzman, K. H. Britten, and W. T. Newsome, Cortical microstimulation influences perceptual judgements of motion direction, *Nature*, **346**, 174-177 (1990).
8. J. M. Simmons and C. R. Gallistel, Saturation of subjective reward magnitude as a function of current and pulse frequency, *Behavioral Neuroscience*, **108**, 151-160 (1994).

Figure Captions

Figure 1. Competition between electrical and gustatory rewards. Panel a depicts the experimental protocol (“Stim” = electrical stimulator). The results in panel b were collected following 17 h of food deprivation. The number of rewards earned per trial is plotted as a function of the number of stimulation pulses, either when the brain stimulation reward (BSR) was presented alone (open triangles) or in competition with a sucrose reward (filled symbols). The results in panels c and d were collected following sodium depletion caused by the administration of a diuretic. Panel c is analogous to panel b with the exception that a saline reward was substituted for the sucrose reward. Panel D shows reward-number curves for BSR alone collected when the subject was salt depleted (open triangles) or salt replete (filled triangles). (“Alt” = alternate; “Std” = standard; “SUC” = sucrose.)

Figure 2. Summation between rewarding effects of lateral hypothalamic (LH) stimulation and sucrose. Panel a depicts the experimental protocol (“Stim” = electrical stimulator). Panel b plots the number of rewards earned as a function of the number of stimulation pulses when the brain stimulation reward (BSR) alternate (filled triangles) competed against a compound standard (filled circles), consisting of intraoral sucrose and an equally-preferred train of LH stimulation. The results obtained with a standard consisting of sucrose alone (from Fig. 1b) are presented for comparison (fine lines). (“Comp” = results of competition test; “Sum” = results of summation test; “Alt” = alternate; “Std” = standard; “SUC” = sucrose.)

Figure 3. Modulation of gustatory reward, but not brain stimulation reward (BSR), by postingestive feedback. Panel a depicts the experimental protocol (“Stim” = electrical stimulator). Reward-number functions obtained for brain stimulation alone before (open triangles) and after (inverted filled triangles) consumption of a large volume (~30 ml) of 1 M sucrose are plotted in panel b. Panel c shows the number of compound (open diamonds) and brain stimulation (open triangles) rewards earned as a function of time when both rewards were

held constant and the gastric cannula was open. The gastric cannula was closed during collection of the results in panel d (filled diamonds: compound reward; filled triangles: BSR alone).

("SUC" = sucrose.)

Figure 4. Two schemes for combining gustatory reward and brain stimulation reward (BSR) (adapted from Conover and Shizgal³). Post-ingestive feedback and homeostatic error signals are depicted as modulating the evaluation of gustatory stimuli (left panel), but not BSR (represented by action potentials in the neurons in the right panel). In the series variant of the model (upper right), the directly activated neurons subserving BSR participate in relaying the result of the gustatory evaluation to the choice mechanism (represented by the pan balance). In the convergence variant (lower right), the directly-activated neurons subserving BSR do not receive a gustatory input, but their output is combined with the result of the gustatory evaluation and relayed to the choice mechanism. We have proposed¹ an experimental strategy for distinguishing between these two variants of the model.



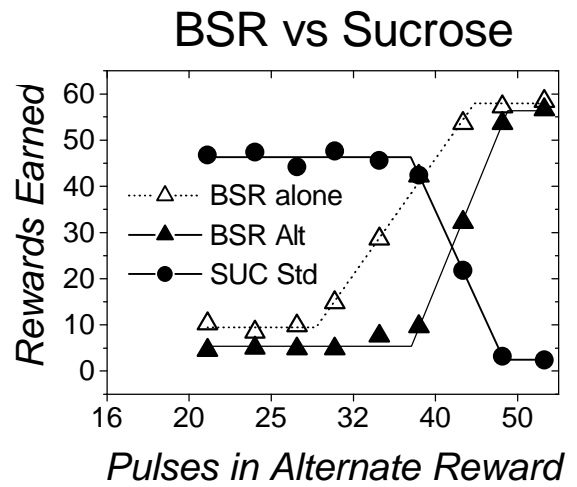



Figure 1






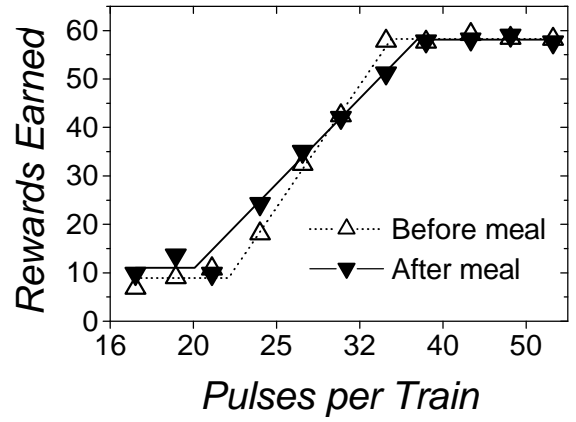
Figure 2

Modulation Protocol



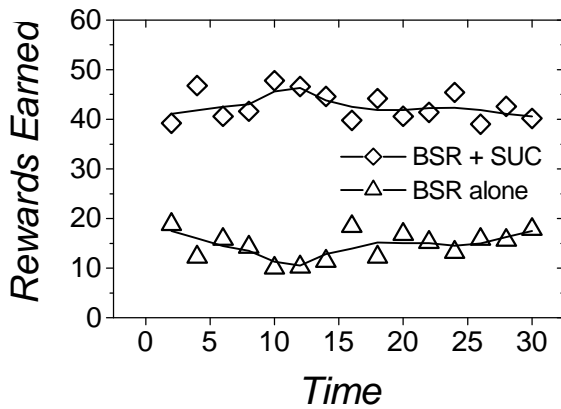
A

BSR Alone




B

Gastric Cannula Open



C

Gastric Cannula Closed



D

Figure 3

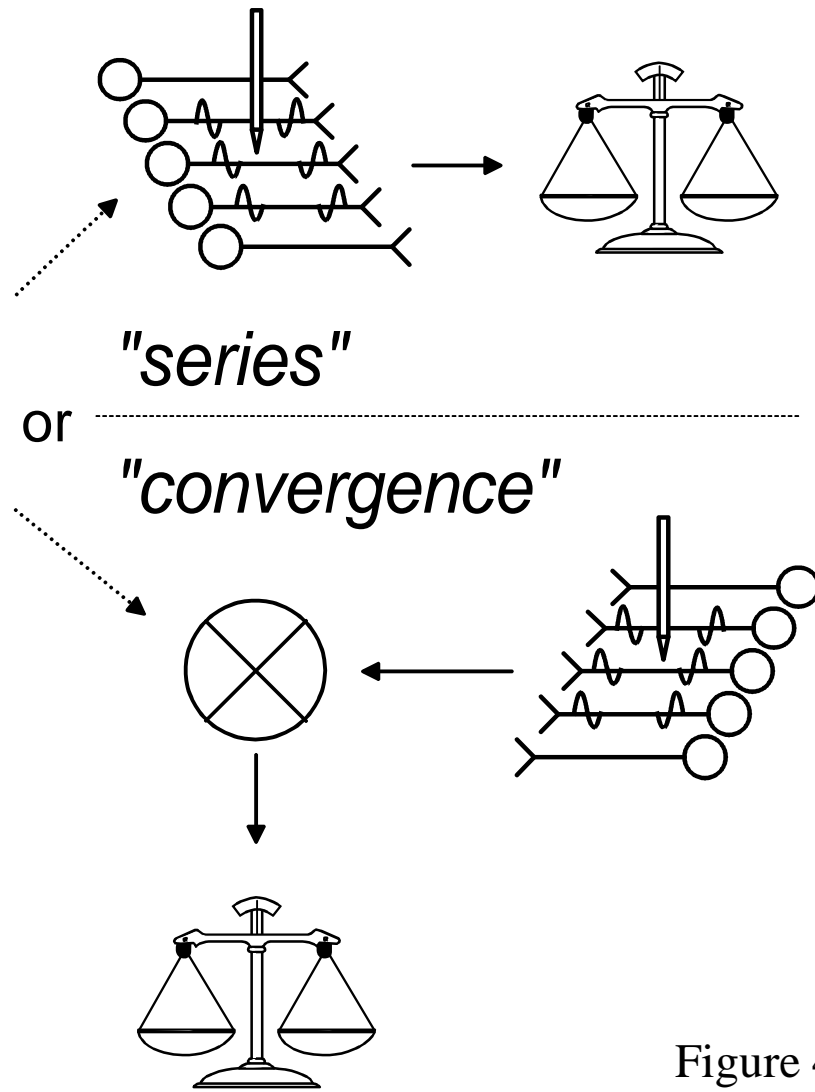
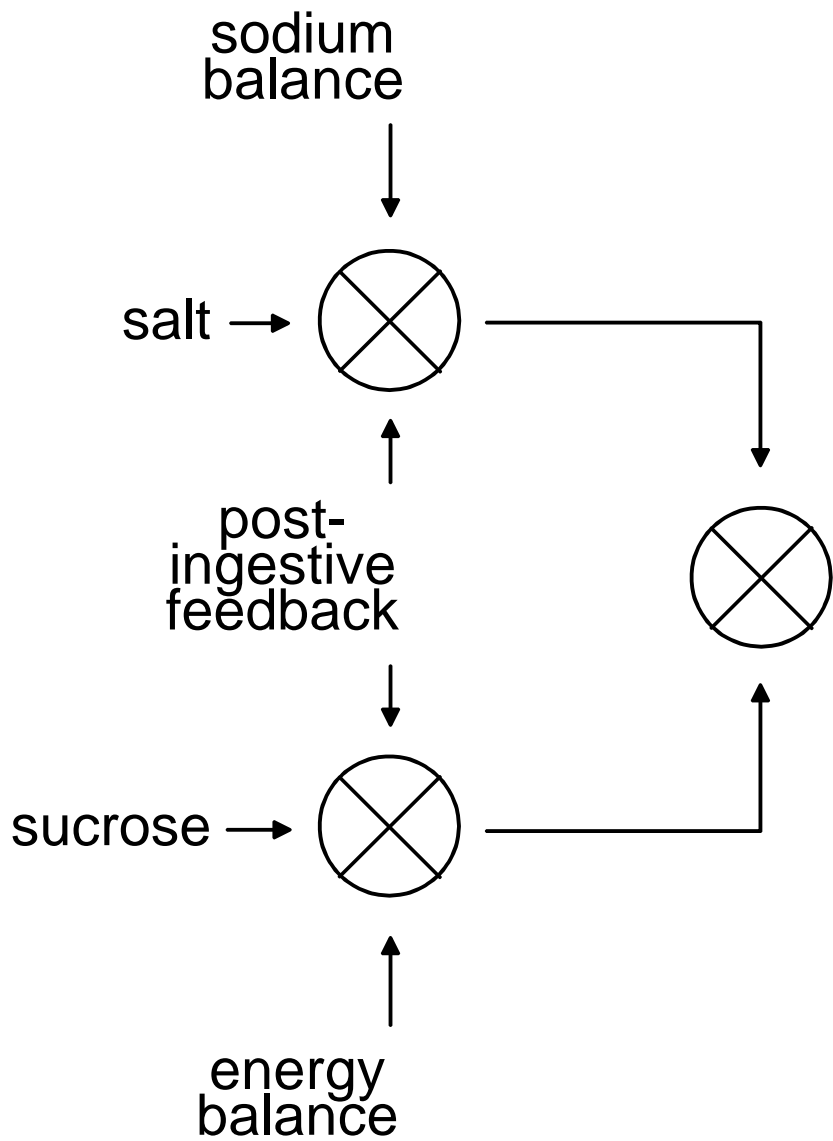


Figure 4